

CLAIMS

What is claimed is:

1. A method comprising:

packing a cache line of each of a plurality of read data returns into one or more packets;

splitting each of the one or more packets into a plurality of flits; and

interleaving the plurality of flits of each of the plurality of read data returns.

2. The method of claim 1, further comprising sending the interleaved flits via a packetized interconnect.

3. The method of claim 1, further comprising receiving the plurality of read data returns from a plurality of memory channels in a substantially overlapped manner.

4. The method of claim 3, wherein a critical chunk of an oldest read data return in a queue is sent in one or more first flits and a critical chunk of a second oldest read data return in the queue is sent in one or more second flits.

5. The method of claim 3, further comprising:

adding a header to each of the plurality of read data returns; and
sending the header before each of the plurality of read data returns.

6. An apparatus comprising:
 - a first buffer to temporarily hold a first cache line of a first read data return;
 - a second buffer to temporarily hold a second cache line of a second read data return; and
 - a multiplexer coupled to the first and second buffers to interleave a first and a second pluralities of flits of the first and second cache lines, respectively.
7. The apparatus of claim 6, further comprising an interface to output the interleaved flits in two packets.
8. The apparatus of claim 7, wherein the multiplexer time-multiplexes the first and the second pluralities of flits in a plurality of time slots to interleave the first and second pluralities of flits.
9. The apparatus of claim 8, wherein the multiplexer dynamically time-multiplexes the first and the second pluralities of flits.
10. The apparatus of claim 8, wherein the multiplexer statically time-multiplexes the first and the second pluralities of flits.
11. The apparatus of claim 7, wherein the interleaved flits are sent via a packetized interconnect to a processor.

12. The apparatus of claim 11, wherein a critical chunk of the first read data return is sent in one or more flits of the first plurality of flits and a critical chunk of the second read data return is sent in one or more flits of the second plurality of flits.

13. The apparatus of claim 6, wherein a header is added to each of the first and second cache lines.

14. The apparatus of claim 11, wherein the header is sent after the corresponding read data return starts arriving at one of the first and the second buffers.

15. The apparatus of claim 11, wherein the header is sent before the corresponding read data return starts arriving at one of the first and the second buffers.

16. The apparatus of claim 6, wherein the first and second read data returns arrive from a first memory channel and a second memory channel, respectively, in a substantially overlapped manner.

17. The apparatus of claim 6, further comprising:

a third buffer, coupled to the multiplexer, to temporarily hold a third cache line of a third read data return, wherein the multiplexer interleaves a third plurality of flits of the third cache line with the first and second pluralities of flits.

18. The apparatus of claim 17, further comprising:

a fourth buffer, coupled to the multiplexer, to temporarily hold a fourth cache line of a fourth read data return, wherein the multiplexer interleaves a fourth plurality of flits of the fourth cache line with the first, the second, and the third pluralities of flits.

19. A system comprising:

a first plurality of dynamic random access memory (“DRAM”) devices;
a second plurality of DRAM devices;
a DRAM channel coupled to the first plurality of DRAM devices;
a second DRAM channel coupled to the second plurality of DRAM devices; and
a memory controller coupled to the first and second DRAM channels, the memory controller including
a first buffer to temporarily hold a first cache line of a first read data return from the first DRAM channel;
a second buffer to temporarily hold a second cache line of a second read data return from the second DRAM channel; and
a multiplexer coupled to the first and second buffers to interleave flits of the first and second cache lines.

20. The system of claim 19, wherein the memory controller sends the interleaved flits in two packets.

21. The system of claim 20, wherein the multiplexer time-multiplexes the first and the second pluralities of flits in a plurality of time slots to interleave the first and second pluralities of flits.
22. The system of claim 21, wherein the multiplexer dynamically time-multiplexes the first and the second pluralities of flits.
23. The system of claim 21, wherein the multiplexer statically time-multiplexes the first and the second pluralities of flits.
24. The system of claim 20, further comprising a packetized interconnect coupled to the memory controller to send the interleaved flits.
25. The system of claim 19, wherein a critical chunk of each of the first and second read data returns is sent in one or more flits.
26. The system of claim 19, wherein the memory controller receives the first and second read data returns in a substantially overlapped manner.
27. The system of claim 19, further comprising a processor coupled to the memory controller to receive the interleaved flits of the first and second cache lines.

28. The system of claim 27, wherein the processor comprises a demultiplexer to separate the flits received.

29. The system of claim 19, further comprising:

- a third plurality of DRAM devices; and
- a third DRAM channel coupled to the third plurality of DRAM devices and the memory controller, wherein the memory controller further includes:
 - a third buffer, coupled to the multiplexer, to temporarily hold a third cache line of a third read data return from the third DRAM channel, wherein the multiplexer interleaves a third plurality of flits of the third cache line with the first and second pluralities of flits.

30. The system of claim 29, further comprising:

- a fourth plurality of DRAM devices; and
- a fourth DRAM channel coupled to the fourth plurality of DRAM devices and the memory controller, wherein the memory controller further includes:
 - a fourth buffer, coupled to the multiplexer, to temporarily hold a fourth cache line of a fourth read data return from the fourth DRAM channel, wherein the multiplexer interleaves a fourth plurality of flits of the fourth cache line with the first, the second, and the third pluralities of flits.

31. A method comprising:

interleaving a plurality of flits containing a critical chunk of each of a first and a second cache lines corresponding to a first and a second read data returns, respectively;

sending the interleaved flits; and

sending a second plurality of flits containing the first cache line's non-critical chunks after the interleaved flits are sent.

32. The method of claim 31, further comprising:

sending a third plurality of flits containing the second cache line's non-critical chunks after the second plurality of flits are sent.

33. The method of claim 32, wherein the first and second read data returns are from a first and a second memory channels, respectively.

34. The method of claim 31, further comprising:

receiving the first and the second read data returns in a substantially overlapped manner.

35. A method comprising:

interleaving a plurality of flits containing a critical chunk of each of a first, a second, and a third cache lines corresponding to a first, a second, and a third read data returns, respectively;

sending the interleaved flits; and

sending a second plurality of flits containing the first cache line's non-critical chunks after the interleaved flits are sent.

36. The method of claim 35, further comprising:

sending a third plurality of flits containing the second cache line's non-critical chunks after the second plurality of flits are sent; and
sending a fourth plurality of flits containing the third cache line's non-critical chunks after the third plurality of flits are sent.

37. The method of claim 36, wherein the first, the second, and the third read data returns are from a first, a second, and a third memory channels, respectively.

38. The method of claim 35, further comprising:

receiving the first, the second, and the third read data returns in a substantially overlapped manner.

39. A method comprising:

interleaving a plurality of flits containing a critical chunk of each of a first, a second, a third, and a fourth cache lines corresponding to a first, a second, a third and a fourth read data returns, respectively;

sending the interleaved flits; and

sending a second plurality of flits containing the first cache line's non-critical chunks after the interleaved flits are sent.

40. The method of claim 39, further comprising:

 sending a third plurality of flits containing the second cache line's non-critical chunks after the second plurality of flits are sent;

 sending a fourth plurality of flits containing the third cache line's non-critical chunks after the third plurality of flits are sent; and

 sending a fifth plurality of flits containing the fourth cache line's non-critical chunks after the fourth plurality of flits are sent.

41. The method of claim 40, wherein the first, the second, the third, and the fourth read data returns are from a first, a second, a third, and a fourth memory channels, respectively.

42. The method of claim 39, further comprising:

 receiving the first, the second, the third, and the fourth read data returns in a substantially overlapped manner.

43. A method comprising:

 checking whether a buffer holds a critical chunk of a cache line of an oldest read return in a queue;

 sending the critical chunk if the buffer holds the critical chunk;

 checking whether a predetermined number of non-critical chunks of the cache line have accumulated in the buffer after the critical chunk is sent; and

sending the non-critical chunks if the predetermined number of non-critical chunks have accumulated in the buffer.

44. The method of claim 43, further comprising:

removing the oldest read return from the queue after sending the non-critical chunks.

45. The method of claim 44, wherein the critical chunk and the non-critical chunks are sent via a packetized interconnect.